

3D SURFACE MAPPING USING A SEMI-AUTONOMOUS ROVER: A PLANETARY ANALOG FIELD EXPERIMENT

Rehman S. Merali¹, Chi Hay Tong¹, Jonathan Gammell¹, Joseph Bakambu², Erick Dupuis³, and Timothy D. Barfoot¹

¹University of Toronto Institute for Aerospace Studies, Toronto, Ontario, Canada, M3H 5T6, {rehman.merali, chihay.tong, jon.gammell, tim.barfoot}@utoronto.ca

²MDA Space Missions, Brampton, Ontario, Canada, L6S 4J3, joseph.bakambu@mdacorporation.com

³Canadian Space Agency, Saint-Hubert, Québec, Canada, J3Y 8Y9, erick.dupuis@asc-csa.gc.ca

ABSTRACT

This paper describes a proposed operational architecture for a planetary worksite mapping mission concept. To map three-dimensional (3D) planetary terrain, we propose to use a rover equipped with a laser rangefinder, and employ a *stop-scan-go* approach with a human-in-the-loop.

In the operational cycle, the rover collects locally consistent 3D range data while stationary. The range data are coupled with visual odometry to estimate the rover pose at each scan and create a consistent 3D map. The 3D map is then used to evaluate candidate next-best views (NBV). The operator selects a NBV with the aid of three evaluation criteria and the rover autonomously travels to the NBV using a network of reusable paths (NRP). Finally, the rover collects another 3D scan and the cycle repeats.

This mission concept was validated through hardware experiments on the CSA's Mars Emulation Terrain (MET), which measures 60m × 120m and includes inclines, rocks, cliffs and a 5.5m-diameter crater.

1. INTRODUCTION

1.1. Motivation

In May 2007, representatives from 14 international space organizations devised the Global Exploration Strategy [1]. The strategy states that robotic exploration is a necessary precursor to human exploration of space. More recently, Tompkins et al. [14] identified water in the Cabeus crater near the South pole of the Moon. It is therefore likely that lunar exploration will begin in a crater at the South pole. The South pole has also been deemed the most probable location for a permanent base station on the Moon [5].

Before a permanent base station is erected on the lunar surface, the site must be accurately mapped in order to further plan the base station. The base station and its surrounding area, where humans and rovers will repeatedly visit the same locations, is defined as a *worksite*.



Figure 1: Operational cycle for 3D site mapping. The rover collects 3D range data while stopped. A batch SLAM algorithm computes the rover pose and creates a consistent 3D map. The operator selects the next-best view (NBV) based on three computed metrics. The rover autonomously drives to the NBV and the cycle repeats.

Planning for both human and robot operations will likely require a consistent three-dimensional (3D) map of the worksite.

While planetary maps have been produced using orbital imagery, these maps are insufficient for some surface operations due to their low resolution. Conversely, a rover on the planetary surface would be able to use a more precise range sensor because the distances being measured would be much less. Given the communication lag from a planetary surface to Earth and potential difficulty communicating from inside a crater, an autonomous rover has distinct benefits over a teleoperated one. Thus, this paper presents a site mapping capability for a rover equipped with a laser rangefinder (as depicted in Figure 1 and Figure 2) to produce consistent surface maps. Lasers are suitable for a planetary environment because they work consistently across a wide variety of lighting conditions including in the dark (ice on the Moon is likely in permanently shadowed regions), have long range, and have low



Figure 2: Rover traversing Mars Emulation Terrain (MET). The rover uses stereo vision for navigation and obstacle avoidance, and a 3D laser rangefinder for mapping and localization.

measurement error. However, collecting a set of 3D range measurements is a slow process. Therefore, the system outlined in the paper employs a *stop-scan-go* approach. For operational safety, a human operator is incorporated in the operational cycle presented in Figure 1, but extensions are proposed for a fully autonomous system.

1.2. Literature Review

Related work conducted by Fong et al. [7] recently demonstrated that it is feasible to conduct a site survey using a rover-mounted laser rangefinder. Their experiments in a planetary analog environment highlighted the benefit of using rovers for the tedious and repetitive task of site surveying. Fong et al. [6] identified additional mission concepts that could also be facilitated by mapping technology, such as resource prospecting and autonomous reconnaissance. However, the experiments presented in [7] and [6] relied on GPS for localization. Since such an infrastructure does not currently exist for Mars or the Moon, this paper presents a framework suitable for GPS-denied environments.

Wettergreen et al. [17] have also demonstrated autonomous rover traverses in planetary analog terrain. While they used GPS for some experiments, Wettergreen et al. also developed a dark navigation system for GPS-denied environments. They used a downward-facing optical sensor and an inertial measurement unit (IMU) and performed dense data alignment to estimate the rover's pose. This system yielded results of 2-3% error on distance traveled. In contrast, the batch SLAM method presented in this paper seeks to produce a globally consistent map and is thus not subject to the drift error encountered by odometry methods.

The 6DSLAM work by Nüchter et al. [11] first utilized the iterative closest point (ICP) algorithm [3] to compute pairwise alignments between scans, then used a global relaxation technique for refinement. This approach resolves the inconsistencies between pairwise linkages by distributing the error over the entire map. However, the

overall performance is still limited by the quality of the initial guesses because pairwise ICP still remains at the core of the approach. In addition, long loops were not detected because loop closures were indicated by a simple distance criterion [18]. In this paper, we utilize a combination of sparse features and odometry measurements in the batch alignment. This approach is similar to GraphSLAM [13], but extended to 3D, and with a different rotation linearization method. The use of sparse features provides implicit loop closure detection, whereas the inclusion of odometry adds flexibility in the distance traveled between scans.

1.3. Outline

The overall operational cycle that will be detailed throughout this paper is shown in Figure 1. To map the 3D terrain using a rover, the rover collects range data while stopped, then uses this data to localize itself and plan its next scan location. Collecting 3D range data while stationary simplifies the scan matching [15]. Visual odometry with a stereo camera is used for high-rate localization during rover movement.

After each 3D laser scan, a batch simultaneous localization and mapping (SLAM) algorithm uses peaks in the terrain to estimate the relative pose of the rover at each scan location. The human operator is shown the amalgamated terrain data and they select candidate locations for the next-best view (NBV). Each candidate NBV is evaluated on three criteria: traversability, localizability and information gain. The operator selects the NBV based on the three criteria, in addition to any waypoints, and the rover plans a path and drives to the NBV autonomously.

A network of reusable paths (NRP) is used to plan and execute the rover's path across the 3D terrain. The NRP software uses stereo vision for navigation and is able to return to any previously driven position with centimeter-level accuracy. Once the rover reaches the NBV, it performs another 3D laser scan, and the operational cycle repeats.

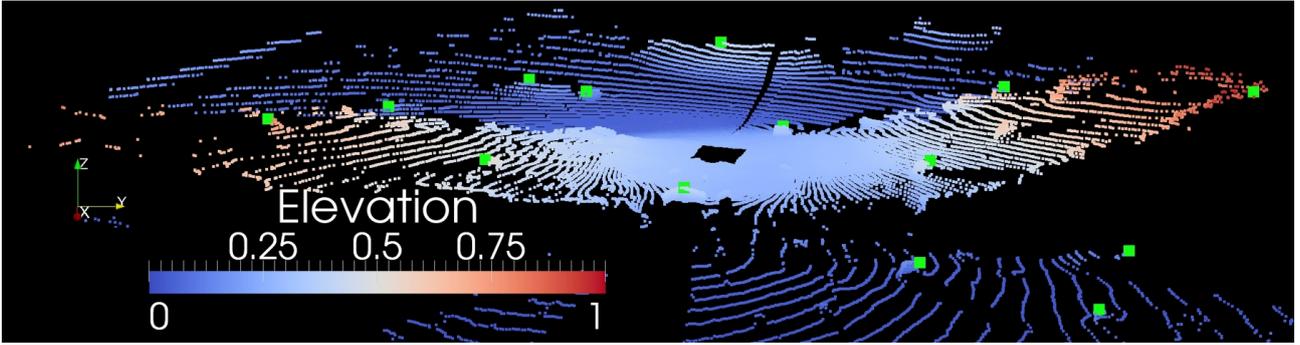


Figure 3: A sample point cloud obtained by a rover equipped with a 3D laser rangefinder in the CSA’s Mars Emulation Terrain (MET), with the world coordinate frame depicted on the left side. The scan was taken from the center of the image, with the rectangular cutout corresponding to the rover’s footprint, and the green squares indicating the detected peak features.

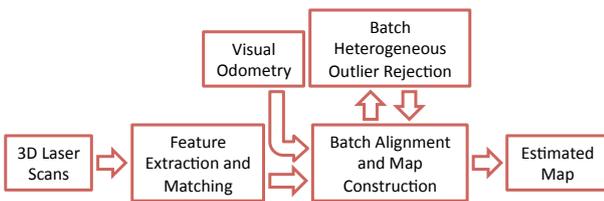


Figure 4: The batch SLAM algorithm used to estimate the rover pose at each scan location. Feature measurements are extracted from the 3D laser scans and coupled with visual odometry data for batch alignment. Batch heterogeneous outlier rejection adds robustness. After convergence, the pose estimates are validated and the scans are overlaid to construct the final 3D map.

The remainder of the paper is organized as follows. Sections 2 to 4 describe the core algorithms developed for the operational cycle. Specifically, Section 2 details the batch SLAM algorithm and Section 3 discusses the NBV evaluation and illustrates the operator interface. Finally, Section 4 outlines the NRP algorithm. After describing the operational cycle and the various algorithms used, Section 5 highlights the field tests conducted at the Canadian Space Agency’s (CSA) Mars Emulation Terrain (MET) in Québec, Canada. The MET measures $60\text{m} \times 120\text{m}$ and includes inclines, rocks, cliffs and a 5.5m-diameter crater. Figure 2 depicts our rover exploring the MET. Through three days of field trials, we tested various operational procedures, culminating in a 25-scan exploratory traverse conducted over the course of a day. Finally, Section 6 concludes with lessons learned and provides suggested avenues for future improvement.

2. BATCH SIMULTANEOUS LOCALIZATION AND MAPPING

The process of collecting 3D range data with a 2D laser rangefinder (LRF) on a pan-tilt unit (PTU) takes several minutes. Thus, acquiring this data while the rover is in motion would introduce significant motion distortion in the data and require the accurate computation of the rover’s pose for the entire traverse. In other words, the

stop-scan-go approach is known to produce more accurate maps [15] because the range data at each 3D scan is locally consistent and only the rover’s pose at each scan location must be computed to resolve the scans in a consistent map.

After the rover has collected a set of locally consistent 3D range data, the batch simultaneous localization and mapping (SLAM) algorithm seeks to align these scans to obtain the six-degree-of-freedom pose of the rover at each scan location, relative to its initial pose. Figure 4 illustrates the batch SLAM algorithm and the details are presented by Tong et al. [15]; however, this section will provide an overview of the algorithm.

Given the 3D range data, the algorithm first extracts distinctive features from each scan. In our implementation, peaks in the terrain are chosen to serve as the distinctive interest points due to their ease of extraction and visibility at long ranges. Though the choice of peaks as features limits the feature extraction and matching approach presented in this section to 2.5D environments, these features produced good results for the planetary mapping scenario. In fully 3D scenarios, an alternative approach may need to be used. Figure 3 shows the peak features detected in a 3D point cloud gathered on the MET.

Matching the peak features between scans is a multistep process. For each pair of scans, constellations of features are matched using the data-aligned rigidity-constrained exhaustive search (DARCES) algorithm [4] to generate hypotheses. This constellation-matching process can be viewed as seeking sets of similar triangles between scans. Unfortunately, due to the large uncertainties in the feature measurements, the DARCES algorithm tends to produce a large number of hypotheses. To reduce this number to only the correct hypotheses, a quality metric is computed by overlaying the dense data for each match. Finally, the data association process is concluded by resolving the valid hypotheses back into individual feature associations.

In addition to the natural features obtained from the range data, we utilize the between-scan visual odometry mea-

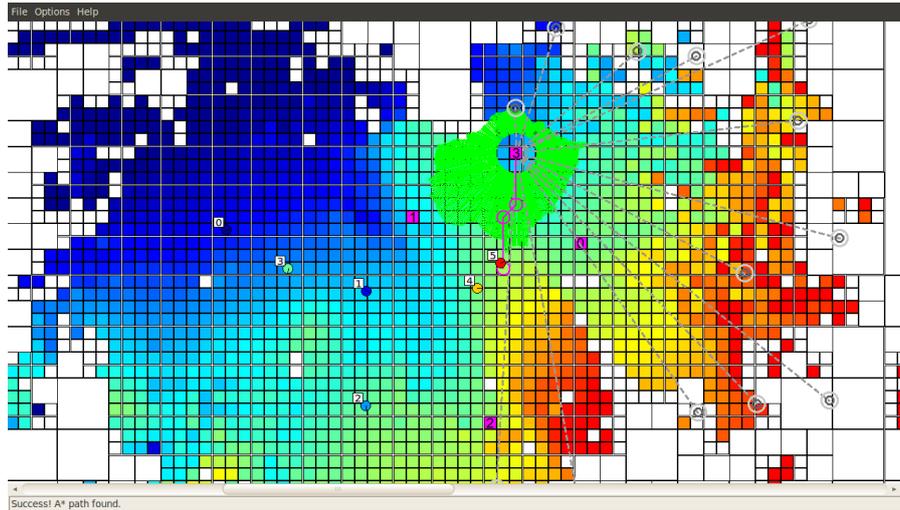


Figure 5: The NBV software displays the merged 3D range data as a 2.5D elevation map where dark blue is the lowest and red is the highest elevation. All rover scan locations are shown and enumerated. Landmarks from the batch SLAM algorithm are shown two concentric gray circles. The user can select any point on the map and the software will compute the shortest traversable path (shown as a pink line), the amount of expected information gain (shown as green lines), and the number of visible landmarks (shown as gray lines).

surements obtained during navigation. This provides flexibility in the rover traverse, because the visual odometry estimates pose-to-pose scan relations in regions of feature scarcity, and situations without any scan overlap. Since these measurements are already utilized during navigation between scan stops, it is straightforward to store them for later use. If these measurements are appropriately incorporated into the batch SLAM algorithm, the presence of more information benefits alignment accuracy.

To maintain global consistency of the estimate, we use a batch approach to the alignment problem. By formulating the SLAM problem as an optimization problem, we seek the optimal estimate of all rover poses and feature locations that best matches the measurement data available. For robustness, the alignment algorithm is augmented with heterogeneous outlier rejection to address the possibility of outliers in either measurement type [16]. After convergence, a post-alignment automatic verification check is performed for each pose estimate to ensure that the alignment produced is valid for inclusion into the final map. This check is computed using the confidence measure obtained from the SLAM algorithm.

The outputs of the batch SLAM algorithm are the pose of the rover at each scan location, the landmark locations, and the merged globally consistent point cloud. These outputs are presented to the operator, and provided to the next-best-view software.

3. NEXT-BEST VIEW

To determine the next scan location, we created a tool to assist the operator. Once the batch SLAM algorithm has computed the rover pose at each scan location thus far, the dense laser data can be combined to produce a consistent

global map. Our software combines the dense laser data in three formats. The first is a point cloud (depicted in Figure 7 and Figure 9), which is displayed to the operator to visually verify the output of the batch SLAM algorithm. Later, this point-cloud view is used to add waypoints en route to the next-best view (NBV), but first the operator must select a NBV. Second, the dense laser data is combined to create a 2.5D elevation map of the terrain. The elevation map is illustrated in Figure 5 and the graphical user interface (GUI) allows the operator to select candidate NBV locations; each of which is evaluated on the three criteria of traversability, localizability and information gain. Traversability can be computed on the 2.5D elevation map, but the other two metrics require a full 3D representation of the laser range data because they require knowledge of the unoccupied space as well as the occupied space. Therefore, all of the positive and negative measurements¹ are combined in a 3D occupancy grid (OG) representation. The 3D OG is written to file and may be viewed at any time. However, viewing the OG is not required during the operational cycle because the localizability and information gain metrics are computed on the OG and displayed to the operator on the GUI.

The GUI shown in Figure 5 illustrates that the 2.5D elevation map is augmented with previous rover scan locations and that when the operator selects a candidate NBV it is shown as a numbered pink square and further augmented by the three metrics. First, the software computes the A* path [9] on the elevation map, where the rover is constrained in the step-size and slope that it can traverse. The traversable step-size, slope and A* grid resolution are all

¹Each range measurement traverses many cells (or voxels) before being reflected. We define the measurement of unoccupied cells as *negative* measurements and the single measurement of an occupied cell as a *positive* measurement.

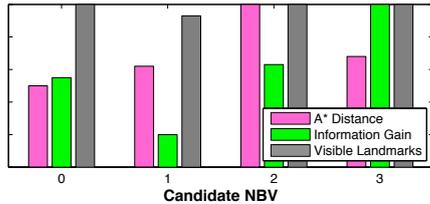


Figure 6: The user is able to compare the selected candidate NBVs based on the three criteria: (i) travel distance, (ii) expected information gain, and (iii) number of visible landmarks.

parameters in the software than can be modified by the user. The A* path is displayed on the GUI as a pink path from the current rover location to the selected candidate NBV. Second, the software uses the 3D OG to trace a ray from the selected NBV to each landmark that was previously identified by the batch SLAM algorithm. Each landmark is shown on the GUI as two concentric gray circles. If a landmark is visible from the candidate NBV, a gray line is drawn between the two. The number of visible landmarks is the localizability metric. Finally, the NBV software computes the expected information gain at each candidate NBV. The expected information gain is approximated by tracing a ray in the OG in every 3D direction (resolution is user-defined) from the candidate NBV and summing the entropy of every visible cell. Each ray is *expected* to terminate when it reaches an occupied cell in the OG or it reaches the maximum range of the sensor. A similar heuristic has been successfully applied in the past; Makarenko et al. [10] used a similar heuristic to approximate expected information gain in 2D. The 3D information gain metric is projected into the 2D plane of the elevation map and displayed on the GUI as green lines originating from the candidate NBV. The length of each green line in the GUI is proportional to the amount of expected information gain in that direction.

All three metrics for each selected candidate NBV are displayed on a bar graph as shown in Figure 6. The operator may again select any previously chosen candidate NBV to visualize the three metrics on the GUI, or select new locations. Once the operator has chosen which of the candidate NBVs is best, this goal location is passed to the Ground Station software depicted in Figure 7 where the user may optionally select waypoints en route to the next-best view. The optional waypoints are locations that the rover will traverse, but not stop and collect a 3D scan. The operator may opt to use waypoints to force the rover along a specific route to the NBV.

4. AUTONOMOUS 3D NAVIGATION

The rover autonomously navigates to the selected next-best view, visiting any optional waypoints along the way, using the network of reusable paths (NRP) as described by Stenning and Barfoot [12]. Physical obstacles are detected by a stereo camera, while an inclinometer is used to detect rollover hazards. The stereo camera also provides visual odometry and localization against a map

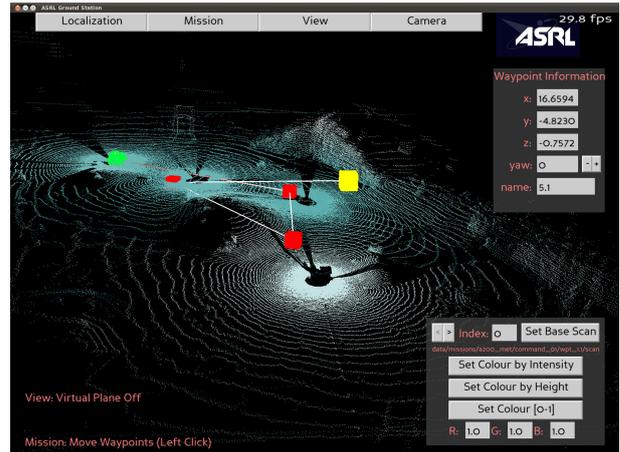


Figure 7: Ground Station software illustrating the map thus far as a 3D point cloud. Four scans have been taken (large green and red boxes) and the fifth is being planned (large yellow box).

of relative pose transformations identified by visual data. Barfoot et al. [2] explain that one of the advantages of the relative pose method is that the rover's pose never needs to be resolved to a single privileged coordinate frame (unlike batch SLAM).

NRP is a planning and navigation algorithm capable of having the rover autonomously drive to its goal location using only a stereo camera. NRP is built upon the visual teach and repeat (VT&R) framework developed by Furgale and Barfoot [8] for navigation in partially-known terrain. Starting with only the rover's current location and the existing paths, NRP uses a rapidly-exploring random tree (RRT) to plan a route from the existing network to the goal in two discrete parts: (i) a retraversal of existing paths to return to a previously-driven pose (the departure pose) and (ii) a kinematically-feasible exploration path from the departure pose to the specified goal, avoiding known obstacles. The NRP algorithm treats each waypoint as a goal.

The rover navigates to the departure pose using the *repeat* mode of the VT&R framework and then switches to *teach* mode to execute the proposed exploration path such that it is simultaneously added to the network. The rover estimates its relative motion while driving the exploration path using visual odometry from the stereo camera. If the rover senses an obstacle along the proposed path, it stops and NRP replans for a new route to the goal that uses the new information collected. This continues until the rover successfully arrives at the specified goal. As a result of the underlying VT&R framework, the rover arrives at the goal with the same positional uncertainty (with respect to the first rover pose) as if it had driven that route directly from the first rover pose to the goal.

Once the rover reaches the next-best-view location, it stops and waits for operator confirmation before collecting another set of 3D laser range data and restarting the operational cycle. The operator can end the cycle at any

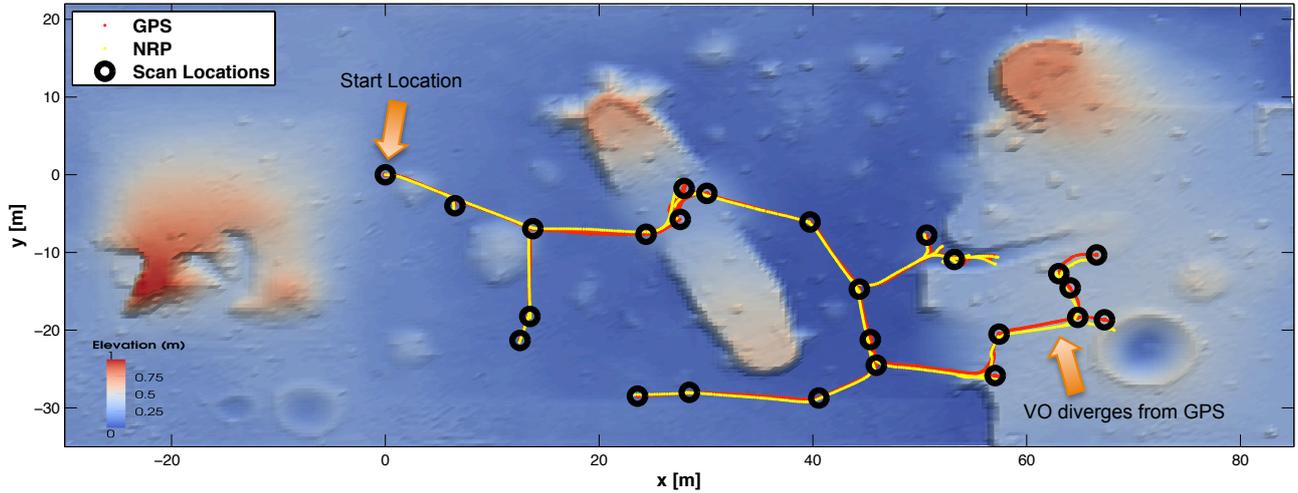


Figure 8: The path traversed during the 25 scan experiment on the MET. Note that the VO localization estimate (used for NRP) diverges from the GPS locations; especially as the distance from the start increases. Notice that the rover retraces its path before branching off from the network to explore new terrain. The rover is able to return to any location on the NRP map with centimeter-level accuracy.

time and command the rover to autonomously return to its starting location. In order to increase system robustness, NRP runs independently of the mapping and NBV software; however, the planning algorithm would benefit from the long-range laser data if the two could be integrated.

For example, Figure 8 shows the 25-scan traverse from the hardware experiments, overlaid on a topological map of the MET. Note that the NRP localization estimate diverges from the GPS ground truth; especially as the distance from the start increases. The 25 scan locations are indicated on the figure (true GPS locations) as large circles. The batch SLAM algorithm is used to ensure that the final 3D map is consistent, as opposed to using the VO estimate of the scan locations. As with all odometry systems, the VO localization error grows with distance from the start. However, the VT&R framework only uses relative transformations to move between positions on the network; therefore, the drift only affects the goal definition and not the rover’s ability to navigate along the network.

5. HARDWARE EXPERIMENTS

As depicted in Figure 2, the entire surface mapping system was tested on physical hardware at the Canadian Space Agency’s (CSA) Mars Emulation Terrain (MET). The rover used was a Husky A200 from Clearpath Robotics; it is a four-wheeled electric vehicle, approximately one meter in length. We equipped the rover with a SICK LMS laser rangefinder mounted on a pan-tilt unit. We also used a Point Grey Bumblebee XB3 stereo camera for NRP and visual odometry. We used RTK-GPS for ground truth localization, but it was not used in the operational cycle. A single laptop computer on the robot computed visual odometry, NRP, and collected laser data. The batch SLAM, Ground Station and NBV software ran off-board at the operator’s command station.

The MET is a $60\text{m} \times 120\text{m}$ area designed to be a planetary analog environment. In addition to inclines, rocks, and cliffs, the MET has different granularity sand and two large craters. Field trials were conducted over three days, during which we tested various operational procedures, culminating in a 25-scan exploratory traverse.

Figure 8 shows the path driven by the rover during the 25-scan traverse. Note that rover retraces its path along the network before branching from the network to explore new terrain. Figure 9 depicts the SLAM solved point cloud of the 25-scan traverse in the Ground Station software. Figure 10 shows an aerial photograph of the terrain for reference. Finally, Figure 11 illustrates the map created from the 25-scan traverse of the MET. Figure 11b is a 3D estimate of the MET given the batch SLAM estimate of the 25 scan locations and Figure 11a is a similar estimate using post-processed GPS data for the 25 scan locations. Discrepancies between the estimated maps and the aerial photograph (particularly at the top of the image) are explained by a lack of range data. Figure 8 shows the scan locations of the 25 scans; areas of the terrain that were not scanned correspond to a poor estimates of the terrain in Figure 11. Comparing Figure 11b to Figure 11a in the well-mapped areas shows that the estimate is consistent; the major ridges, craters, and flat terrain are accurately depicted.

6. CONCLUSION AND LESSONS LEARNED

The surface mapping system was deemed quite successful in the field. The rover was able to traverse the planetary analog environment and create a consistent 3D surface map. The final map is represented in three different forms, (i) a 3D point cloud, (ii) a 2.5D elevation map, and (iii) a 3D occupancy grid. We were able to incorporate a human in the loop for operational safety; however, a future extension of the work may be to randomly sample candidate NBVs and select the optimal one based on a



Figure 9: Ground Station software illustrating the map after the 25-scan traverse. The terrain is shown as a 3D point cloud, and overlaid are the scan locations (large green and red boxes) with lines between them indicating the order. The green box is the starting location of the mission. The small green circles on the map are landmarks.



Figure 10: An aerial photograph of the Mars Emulation Terrain (MET).

weighted sum of the three metrics: traversability, localizability, and information gain.

Based on these preliminary experimental results, we believe that the operational cycle described in the paper and depicted in Figure 1 is very promising. We were able to successfully map CSA's MET by emulating a planetary deployment. Possible areas for improvement include tighter integration between the various subsystems. For example, the NRP software could benefit from the long-range laser data to plan paths into areas of the map that are not in the network. However, the laser data cannot be incorporated into the network without first resolving the localization discrepancy between VO and batch SLAM. The NBV software could also be integrated with

the Ground Station software, but the decision was made to keep them separate so each can be used separately on future projects. Furthermore, since the 3D data is represented in three different forms, the operator may benefit from switching between these at will, as opposed to performing specific tasks on each view. The 25-scan traverse of the MET produced a consistent 3D map of the planetary analog environment, but it could be improved by more scans. The mapped areas of the MET resemble the true terrain and we believe that they would be sufficient for planetary worksite operations.

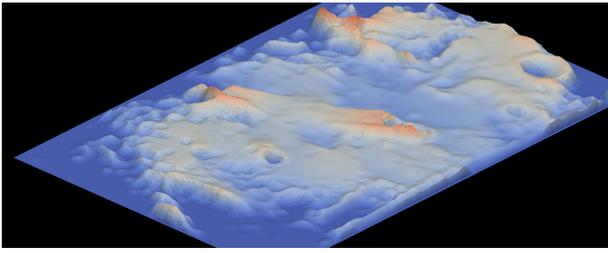
ACKNOWLEDGMENTS

The authors would like to thank Clearpath Robotics for providing the rover platform used in the field experiment. We also thank Braden Stenning for providing the NRP software and Keith Leung for the Ground Station software.

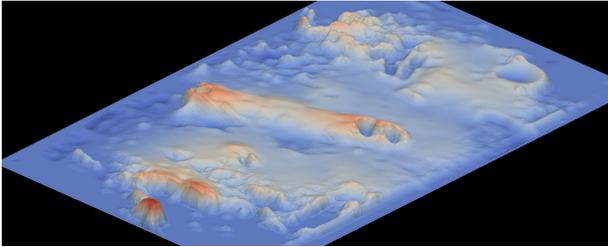
This work was supported by a Natural Sciences and Engineering Research Council of Canada Collaborative Research and Development (NSERC CRD) / Canadian Space Agency Partnership Support Program (CSA PSP) grant in cooperation with MDA Space Missions.

REFERENCES

- [1] ASI (Italy), BNSC (United Kingdom), CNES (France), CNSA (China), CSA (Canada), CSIRO (Australia), DLR (Germany), ESA (European



(a) A 3D render of the MET constructed as a result of ground truth GPS localization.



(b) A 3D render of the MET constructed as a result of the batch SLAM localization estimates.

Figure 11: 3D renders of the MET created from the 25-scan traverse of the terrain.

Space Agency), ISRO (India), JAXA (Japan), KARI (Republic of Korea), NASA (United States of America), NSAU (Ukraine), and Roscosmos (Russia), “The global exploration strategy: The framework for coordination,” Technical report, May 2007.

- [2] T. D. Barfoot, B. Stenning, P. Furgale, and C. McManus, “Exploiting reusable paths in mobile robotics: Benefits and challenges for long-term autonomy,” in *Computer and Robot Vision, Ninth Conference on*, Toronto, Canada, May 2012.
- [3] P. Besl and H. McKay, “A method for registration of 3-d shapes,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 14, no. 2, pp. 239–256, Feb 1992.
- [4] C.-S. Chen, Y.-P. Hung, and J.-B. Cheng, “Ransac-based darces: a new approach to fast automatic registration of partially overlapping range images,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 21, pp. 1229–1234, 1999.
- [5] S. Dale, “Exploration strategy and architecture,” in *Implementing the Vision - 2nd Space Exploration Conference*, 2006.
- [6] T. Fong, M. Bualat, M. Deans, M. Allan, X. Bouyssounouse, M. Broxton, L. Edwards, R. Elphic, L. Flückiger, J. Frank, L. Keely, L. Kobayashi, P. Lee, S. Y. Lee, D. Lees, E. Pacis, E. Park, L. Pedersen, D. Schreckenghost, T. Smith, V. To, and H. Utz, “Field testing of utility robots for lunar surface operations,” in *Proc. AIAA Space 2008*, 2008.
- [7] T. W. Fong, M. Allan, X. Bouyssounouse, M. G. Bualat, M. Deans, L. Edwards, L. Flückiger, L. Keely, S. Y. Lee, D. Lees, V. To, and H. Utz, “Robotic site survey at haughton crater,” in *9th International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, February 2008.
- [8] P. T. Furgale and T. D. Barfoot, “Visual teach and repeat for long-range rover autonomy,” *Journal of Field Robotics*, special issue on “Visual mapping and navigation outdoors”, vol. 27, no. 5, pp. 534–560, 2010.
- [9] P. E. Hart, N. J. Nilsson, and B. Raphael, “A formal basis for the heuristic determination of minimum cost paths,” *IEEE Trans. Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [10] A. Makarenko, S. Williams, F. Bourgault, and H. Durrant-Whyte, “An experiment in integrated exploration,” *IEEE/RSJ International Conference on Intelligent Robots and System*, vol. 1, pp. 534–539, 2002.
- [11] A. Nüchter, K. Lingemann, J. Hertzberg, and H. Surmann, “6d slam - 3d mapping outdoor environments,” *Journal of Field Robotics (JFR), Special Issue on Quantitative Performance Evaluation of Robotic and Intelligent Systems*, vol. 24, pp. 699–722, 2007.
- [12] B. Stenning and T. D. Barfoot, “Path planning on a network of paths,” in *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT, March 2011.
- [13] S. Thrun and M. Montemerlo, “The GraphSLAM algorithm with applications to large-scale mapping of urban structures,” *International Journal on Robotics Research*, vol. 25, no. 5/6, pp. 403–430, 2006.
- [14] P. D. Tompkins, R. Hunt, M. D. D’Ortenzio, J. Strong, K. Galal, J. L. Bresina, D. Foreman, R. Barber, M. Shirley, J. Munger, and E. Drucker, “Flight operations for the LCROSS lunar impactor mission,” in *SpaceOps*, no. AIAA-2010-1986, Huntsville, Alabama, April 2010.
- [15] C. Tong, T. D. Barfoot, and E. Dupuis, “Three-dimensional SLAM for mapping planetary worksite environments,” *Journal of Field Robotics*, special issue on “Space Robotics”, vol. 29, no. 3, pp. 381–412, 2012. [Online]. Available: <http://dx.doi.org/10.1002/rob.21403>
- [16] C. Tong and T. D. Barfoot, “Batch heterogeneous outlier rejection for feature-poor SLAM,” in *Proceedings of the IEEE Conference on Robotics and Automation (ICRA)*, Shanghai, China, 9-13 May 2011, pp. 2630–2637.
- [17] D. Wettergreen, S. J. Moreland, K. Skonieczny, D. Jonak, D. Kohanbash, and J. Teza, “Design and field experimentation of a prototype lunar prospector,” *International Journal of Robotics Research*, vol. 29, no. 12, pp. 1550–1564, October 2010.
- [18] O. Wulf, A. Nüchter, J. Hertzberg, and B. Wagner, “Benchmarking urban six-degree-of-freedom simultaneous localization and mapping,” *Journal of Field Robotics*, vol. 25, no. 3, pp. 148–163, 2008. [Online]. Available: <http://dx.doi.org/10.1002/rob.20234>